arXiv:1704.00352v1 [stat.AP] 2 Apr 2017

Simple Measures of Individual Cluster-Membership Certainty for Hard Partitional Clustering

Dongmeng Liu
Statistics and Actuarial Science, Simon Fraser University
and
Jinko Graham*
Statistics and Actuarial Science, Simon Fraser University

April 4, 2017

Abstract

We propose two posterior-probability-like measures of individual cluster-membership certainty which can be applied to a hard partition of the sample such as that obtained from the Partitioning Around Medoids (PAM) algorithm. One measure extends the individual silhouette widths and the other is obtained directly from the pairwise dissimilarities in the sample. Unlike the classic silhouette, however, the measures behave like probabilities and can be used to investigate an individual's tendency to belong to a cluster. Motivated by an application to a clinical database, we evaluate the performance of both measures in individuals with ambiguous cluster membership, using simulated binary datasets that have been partitioned by the PAM algorithm. For comparison, we also present results from soft clustering algorithms such as soft analysis clustering (FANNY) and two model-based clustering methods. Our proposed measures perform comparably to the posterior-probability estimators from either FANNY or the model-based clustering methods.

Keywords: Cluster-membership certainty, FANNY algorithm, Model-based clustering, Partitioning around medoids algorithm, Silhouette width, Soft clustering

*This work was partially supported by *the Natural Sciences and Engineering Research Council of Canada*. The authors thank Joanna Lubieniecka for helpful discussions about clinical databases.

1 Introduction

Clinical disease registries frequently contain information recorded in the form of categorical variables for each patient. To explore such data, we may wish to cluster the patients into similar groups. For example, we may wish to use patient symptoms at diagnosis to identify groups which respond differently to treatment. One approach to clustering with categorical data is Bayesian profile regression (Molitor et al., 2010), which has the ability to incorporate information on an outcome variable. The profile regression model is fitted to the data by use of a Markov Chain Monte Carlo algorithm, in which the number of clusters and cluster membership changes at each sweep (Liverani et al., 2015) and the co-occurrence of a pair of individuals in the same cluster is tracked. After completion of all sweeps, a similarity matrix is created by averaging the pairwise co-occurrences across the sweeps. Then individuals are assigned to clusters by applying the Partitioning Around Medoids or PAM algorithm (Kaufman and Rosseeuw, 1990) to the resulting dissimilarity matrix.

One limitation of this approach is that so-called "hard" partitional clustering algorithms such as PAM assign individuals to distinct clusters but do not provide a measure of the cluster-membership certainties for each individual. Yet, in many applied settings, cluster-membership certainties are desired to help identify individuals with ambiguous group memberships. One measure of how well an individual belongs to its assigned cluster is the silhouette (Rousseeuw, 1987). Silhouette values range between negative and positive one, with high values indicating that the individual is well matched to its assigned cluster relative to neighbouring clusters. In this note, we propose a simple extension of the silhouette from a single value pertaining to the individual's assigned cluster to a vector of values pertaining to all the clusters in the partition. An attractive feature of the extension is that an individual's values add to one across the clusters and thus provide a posterior-probability-like interpretation. Such an interpretation is helpful for assessing the individual's membership uncertainty after the hard clustering has been performed. We also propose another posterior-probability-like measure of cluster-membership based directly on the dissimilarity matrix and the partition. The performance of the proposed measures is evaluated in a limited simulation study. Both measures behave similarly to posterior probabilities from model-based and fuzzy clustering. For researchers exploring their data with a hard partitional clustering algorithm, the proposed measures therefore offer a straightforward way to augment existing output and obtain posterior-probability-like measures of cluster membership uncertainty. Although our motivation is an application from Bayesian profile regression, the measures can be applied to any pairwise dissimilarity matrix and cluster membership assignment obtained from hard clustering.

2 Proposed Measures

2.1 Silhouette-Based

The silhouette is a widely-used interpretation of how well each individual lies within its assigned cluster (Rousseeuw, 1987). Each individual's silhouette value is defined by comparing the individual's average dissimilarity with others in its assigned cluster to its dis-

similarity with individuals in all other clusters. We extend an individual's silhouette value to a vector, as follows. Given the hard partition, we re-assign the individual to a different cluster holding fixed the other individuals' assignments and compute the corresponding vector of silhouette values for the individual of interest. Since the silhouette values range between -1 and 1, a simple way to make all values positive is to add 1 to every element of the vector. We then add a user-specified exponent, l, to the shifted silhouette values and convert them into probabilities by dividing each element by the sum of all elements.

Let $Z = (z_1, \ldots, z_N)$ denote the cluster membership assignment or partition for the N individuals in the sample and C denote the number of clusters in the partition. For each individual i, we set $z_i = k$ for k in $1 \ldots C$, but leave all remaining elements of Z (for the other individuals) unchanged. Let $s_{ik} \in [-1, 1]$ denote the silhouette value of individual i when the individual is assigned to cluster k. Therefore, each individual i is assigned a vector of silhouette values as $S_i = (s_{i1}, \ldots, s_{iC})$. Let $P_{ik}^{(1)}$ denote the silhouette-based measure of cluster-membership certainty for individual i belonging to cluster k. Then we define $P_{ik}^{(1)}$ as

$$P_{ik}^{(1)} = \frac{(s_{ik} + 1)^l}{\sum_{j=1}^{C} (s_{ij} + 1)^l},$$

where l is a user-specified parameter. Increasing the exponent term l pushes small elements in the shifted silhouette vector closer to zero but has relatively little effect on larger elements that are closer to two. Large values of l should therefore produce crisper clusters and lower misclassification rates.

2.2 Dissimilarity-Based

In addition to the silhouette-based measure, we propose a measure that is based directly on the pairwise dissimilarity matrix. Assume that the pairwise dissimilarity matrix, $\{d_{ij}\}$, between N individuals is given, and has non-negative entries. To ensure dissimilarity measures between 0 and 1, we standardize the (i, j)th entry of the dissimilarity matrix as follows:

$$D_{ij} = \frac{d_{ij} - min\{d_{11}, \ldots, d_{NN}\}}{max\{d_{11}, \ldots, d_{NN}\} - min\{d_{11}, \ldots, d_{NN}\}}.$$

Let $h_{ik} \in [0, 1]$ be the average dissimilarity between individual i and cluster k such that

$$h_{ik} = \left(\sum_{j \neq i : Z_j = k} D_{ij} \right) \Big/ |\{j \neq i : Z_j = k\}|$$

Since h_{ik} is the average dissimilarity between individual i and cluster k, $1 - h_{ik}$ represents the corresponding average similarity. The dissimilarity-based measure of cluster-membership certainty of individual i belonging to cluster j is therefore

$$P_{ik}^{(2)} = \frac{(1 - h_{ik})^v}{(1 - h_{i1})^v + (1 - h_{i2})^v + \ldots + (1 - h_{iC})^v},$$

where v is a user-specified exponent. Increasing v pushes small similarities closer to zero but has relatively little effect on large similarities that are close to one. Large v should therefore produce crisper clusters and lower misclassification rates.

3 Simulation Study

We evaluate the cluster-membership certainty of an individual which is a hybrid of two groups. Group membership is determined by a latent variable, U, taking on two values. From each of the two groups, 20 individuals are simulated with 20 binary features. We assign individuals $1, \ldots, 20$ to group 1 and individuals $21, \ldots, 40$ to group 2. Let $X_{ij} \in \{0, 1\}$ denote the jth feature of individual i, and U_i the latent variable for individual i. The conditional probability of a binary feature being a success is

$$Pr(X_{ij} = 1 \mid U_i) = \frac{exp(t_i + \beta * U_i)}{1 + exp(t_i + \beta * U_i)},$$

where the intercept t_i is selected to ensure that $Pr(X = 1) = \sum_u Pr(X = 1 \mid u) Pr(U_i = u) = 0.5$.

We set $\beta = 1.2$; $U_1 = \ldots = U_{20} = 1$ for group 1; and $U_{21} = \ldots = U_{40} = 4$ for group 2. These settings ensure that the two groups can be easily differentiated on a plot of the first two principal coordinates from a multiple correspondence analysis (see Le Roux and Rouanet, 2004; results not shown). The hybrid individual is then added to the dataset and is assigned half of its features from one group and the other half from the other group. To be specific, features $X_1, \ldots, X_{20}$ are based on $U = 1$ and $X_{21}, \ldots, X_{40}$ are based on $U = 4$. For the 40 non-hybrid individuals, we define clusters 1 and 2 as those assigned more individuals coming from true groups 1 and 2, respectively. An individual is classified correctly if the indices of its true group and cluster assignment agree, and misclassified otherwise. We compute $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$, the two measures of cluster-membership certainty of the hybrid individual for cluster 1 when the number of clusters is fixed to 2.

We consider a number of possible dissimilarity matrices in our simulations of the binary data: Euclidean distance based on the top two principal coordinates from multiple correspondence analysis (see, e.g. Le Roux and Rouanet, 2004), simple matching distance (SMD) (see, e.g. Gower, 2004) and the `PReMiuM` co-occurrence dissimilarity from profile regression described in the introduction. As a benchmark for comparison, we also apply soft-clustering and compute the hybrid's posterior probability of belonging to cluster 1 under (i) LCA applied directly to the discrete data on the features (see, e.g. Lazarsfeld and Henry, 1968 and McCutcheon, 1987), (ii) a Gaussian mixture model (Banfield and Raftery, 1993) applied to the top two principal coordinates obtained from multiple correspondence analysis and (iii) the FANNY algorithm's memberships (Kaufman and Rosseeuw, 1990). We simulate 1000 datasets to obtain the empirical distribution of the hybrid's cluster-membership certainties. Since hybrid individuals have an equal proportion of variables from each contributor group, we expect the distributions of $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$ to be symmetric, with a mean around 0.5.

We also consider the 40 non-hybrid individuals and calculate their "soft" misclassification rate, $M^{(q)}$, where $q = 1, 2$ for the silhouette-based and the dissimilarity-based measure of cluster-membership certainty, respectively. Let g_i denote the true group of individual i; then the soft misclassification rate is $M^{(q)} = \frac{1}{N} \sum_{i=1}^{N} (1 - P_{ig_i}^{(q)})$, for $q = 1$ or 2. The soft misclassification rate weights crisp cluster memberships differently than fuzzy memberships. Specifically, the higher the membership certainty for the true group of an individual, the lower the contribution of that individual to the soft misclassification rate.

We evaluate the proposed measures of cluster-membership certainty using the hard partition assigned by PAM because this algorithm is relevant to our motivating application of profile regression for clinical data. We first investigate how the proposed measures of cluster-membership certainty work with the parition from PAM. We then compare their performance to fuzzy-clustering measures from the FANNY algorithm and to model-based clustering measures from LCA and Gaussian mixture models.

The simulation study is implemented in R. The R package `FactoMineR` (Lê et al., 2008) provides functions for multiple correspondence analysis and data visualization. The PAM and FANNY algorithms are implemented in the R package `cluster`. The implementation of LCA for binary covariates is available in the R package `poLCA` (Linzer and Lewis, 2011; R Core Team, 2012). The Gaussian mixture model is implemented in the R package `mclust` (see Fraley et al., 2012 and Fraley and Raftery, 2002).

4 Results

4.1 Proposed Measures

For both of the proposed measures, the tuning parameter changes the misclassification rate and the distributional shape of the hybrid's cluster-membership certainty (posterior probability). Table 1 shows the effect of tuning the exponent parameter of the two measures for the Euclidean distance matrix, SMD matrix and `PReMiuM` co-occurrence matrix. In the table, the tuning parameters are chosen so that the standard deviations are 0.15, 0.20 and 0.25, and the corresponding misclassification rates are presented. The mean of $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$ is 0.50 regardless of the dissimilarity matrix or the value of tuning parameter, as expected for the hybrid individual. In general, the exponent parameters l and v represent a tradeoff between detecting hybrid individuals and minimizing the soft misclassification rate for the non-hybrid individuals. Specifically, increasing l and v leads to larger variance in $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$ but lower misclassification rates $M^{(1)}$ and $M^{(2)}$. For example, referring to the entry of the silhouette-based measure of the Euclidean distance matrix in the first row of the table, increasing the tuning parameter l from 0.7 to 1.4 increases $\text{sd}(P_{h1}^{(1)})$ from 0.15 to 0.25 and decreases $M^{(1)}$ from 10.04% to 1.60%.

Figure 1 shows an example of how tuning l and v influences the shape of the distribution of $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$ based on the Euclidean distance matrix. For the silhouette-based measure with a large $l = 5$, $P_{h1}^{(1)}$ is more variable, with most values being close to 0 or 1; as l decreases to $l = 0.5$, $P_{h1}^{(1)}$ takes on more extreme values but still has a mode at 0.5. Similarly, for the dissimilarity-based measure, most of the values of $P_{h1}^{(2)}$ are close to either 0 or 1 when $v = 25$ but when $v = 5$ they tend to concentrate around 0.5.

The performance of the two measures depends on the dissimilarity matrices. Referring to Table 1, we see that, for the Euclidean distance matrix, the silhouette-based measure has a larger misclassification rate than the dissimilarity-based measure for a given value of $\text{sd}(P_{h1})$. For example, when $\text{sd} = 0.15$, $M^{(1)} = 10.04\% > M^{(2)} = 4.97\%$. For the SMD matrix, the silhouette-based measure also has a higher misclassification rate than the dissimilarity-based measure for a given value of the standard deviation. However, for the

Table 1: The soft misclassification rates M and $\text{sd}(P_{h1})$ for different values of tuning parameters for the three dissimilarity matrices. The value of the tuning parameter is shown in parenthesis.

		sd=0.15	sd=0.20	sd=0.25
Euclidean	silhouette-based	10.04% (l=0.7)	4.50% (l=1.0)	1.60% (l=1.4)
	dissimilarity-based	4.97% (v=3.8)	2.21% (v=5.2)	0.79% (v=7.1)
SMD	silhouette-based	6.84% (l=3.0)	2.92% (l=4.2)	1.09% (l=5.7)
	dissimilarity-based	5.66% (v=2.3)	2.18% (v=3.3)	0.80% (v=4.5)
`PReMiuM`	silhouette-based	8.18% (l=0.6)	3.09% (l=0.95)	0.70% (l=1.5)
	dissimilarity-based	14.80% (v=0.13)	9.86% (v=0.19)	5.60% (v=0.27)

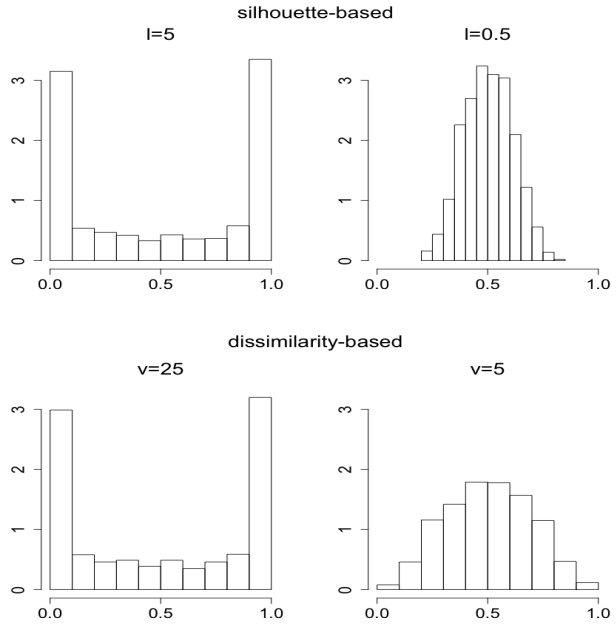

Figure 1: Decreasing l (top) and v (bottom) changes the respective distribution of $P_{h1}^{(1)}$ and $P_{h1}^{(2)}$ from being symmetric with a U-shape to being symmetric with a mode at 0.5, as shown for the Euclidean distance matrix

PReMiuM co-occurrence dissimilarity matrix, the silhouette-based measure has much lower misclassification rates than the dissimilarity-based measure.

4.2 Comparison to Other Clustering Methods

As a bench mark, we compute the estimated posterior probabilities in the sample for the FANNY algorithm, LCA and Gaussian model-based clustering. Table 2 summarizes results for all the clustering methods. In the FANNY algorithm, the user-specified, exponent parameter r affects the misclassification rate for the data. Unlike the parameter l and v of the proposed measures, increasing r creates fuzzier clusters and generates more uncertainty for both non-hybrid and hybrid individuals. Increasing FANNY's exponent parameter r leads to an increase in the overall uncertainty, just like decreasing the exponent parameters v and l of the proposed measures. Therefore, FANNY's parameter r also represents a tradeoff between detecting hybrid individuals and misclassifying non-hybrid individuals. In Table 2, these tuning parameters are set to default values of $v = 1$, $l = 1$ and $r = 2$. FANNY does well at detecting the hybrid individuals, whose cluster-membership certainty measures appear to be normally distributed with a mean around 0.5 (results not shown). However, the larger value of FANNY's misclassification rate for the SMD matrix suggests that decreasing r may be necessary, even though this will increase the variance of the hybrid's measure of cluster certainty.

For LCA and Gaussian model-based clustering, we observe extreme behaviour for the hybrid's cluster-membership certainties: the hybrid individual always has an estimated posterior probability of either zero or one (results not shown), with equiprobable assignment to either extreme. The hybrids' posterior probabilities show little uncertainty, which is incompatible with how the hybrid data were simulated. Although the posterior probability estimators appear to be unbiased (the point estimates are within simulation error of 0.5, at 95% confidence level), the standard deviations reach or almost reach the maximum of 0.5. That is, a hybrid individual is randomly assigned to either cluster with a posterior probability equal to one.

Table 2: Comparison of the results for all the clustering methods

		mean	sd	M
silhouette-based (l=1)	Euclidean	0.50	0.20	4.50%
	SMD	0.50	0.06	28.38%
	PReMiuM	0.50	0.21	2.39%
dissimilarity-based (v=1)	Euclidean	0.50	0.04	29.07%
	SMD	0.50	0.07	21.15%
	PReMiuM	0.50	0.42	0.10%
FANNY (r=2)	Euclidean	0.50	0.16	3.88%
	SMD	0.50	0.04	32.45%
	PReMiuM	0.50	0.17	1.22%
model-based	LCA	0.48	0.50	0.03%
	Gaussian	0.50	0.46	0.02%

5 Discussion

Pairwise dissimilarities between samples reflect the structure present in multivariate data and provide key information for clustering individuals. The silhouette value uses pairwise dissimilarities to measure how well an individual fits to the cluster it has been assigned relative to the other clusters in a hard partition. We have proposed two posterior-probability-like measures of cluster membership certainty, one which extends the classic silhouette and the other based directly on the dissimilarities. These measures can assist with identifying individuals of ambiguous cluster membership after applying a hard clustering algorithm. As they are simple and behave like posterior probabilities, they may be conveniently applied in clinical and bioinformatics settings which use hard partitional clustering to explore data (e.g., Molitor et al., 2010 and Liverani et al., 2015). Motivated by a clinical application with binary data, we evaluate the proposed measures in a small simulation study with selected dissimilarity matrices and find that they perform well.

In our simulations, all the methods provide unbiased measures of the hybrid individual's posterior probability of cluster membership. However, the measures from the model-based methods have a U-shaped distribution and high variance. In contrast, our measures and those from the FANNY algorithm can be tuned to have a distribution with a mode at 0.5. One direct contrast of the proposed measures to the model-based estimates of posterior probability can be seen in the application of PAM and Gaussian model-based clustering to the Euclidean distance matrix of the top principal coordinates from the multiple correspondence analysis. Our measures easily capture the hybrid's ambiguous membership whereas the model-based clustering methods do not.

One feature that our proposed measures and the FANNY algorithm have in common is the way the exponent parameter changes the variance of the hybrid's cluster membership and the misclassification rate in the sample. In general, increasing the exponent parameter leads to an increase of variance of the hybrid's distribution and a decrease of misclassification rate. For the proposed measures, increasing the exponentiation l and v for elements of the similarity matrix pushes the small similarities closer to the lower bound, with relatively little effect on those close to the upper bound, thereby decreasing the uncertainty and creating crisper clusters. The tuning parameters l and v represent a tradeoff between the soft misclassification rate for non-hybrids and the variance of the hybrid's uncertainty measure. We recommend that researchers experiment with tuning l or v to balance this tradeoff. A similar recommendation has been proposed for the tuning parameter r in the FANNY algorithm (see Kaufman and Rosseeuw, 1990).

In our simulations, the proposed measures reflect the hybrid individual's posterior probability of cluster membership as expected, though their misclassification rates are higher than FANNY's. The higher misclassification rates are expected since the proposed measures work from a fixed clustering while FANNY has more flexibility to simultaneously cluster and assign fuzzy memberships.

In summary, our measures are straightforward to implement and worth considering as a way to augment PAM and other hard clustering methods which give no measure of the posterior probabilities of cluster membership for individuals.

References

Jefferey D. Banfield and Adrian E. Raftery. Model-based Gaussian and non-Gaussian clustering. *Biometrics*, 49(3):803–821, 9 1993.

C. Fraley and A. E. Raftery. Model-based clustering, discriminant analysis and density estimation. *Journal of the American Statistical Association*, 97:611–631, 2002.

C. Fraley, A. E. Raftery, T. B. Murphy, and L Scrucca. *mclust Version 4 for R: Normal Mixture Modeling for Model-Based Clustering, Classification, and Density Estimation*, 2012.

J. Gower. *Similarity, dissimilarity and distance, measures of.* Encylcopedia of statistical sciences, 2004.

L. Kaufman and P. J. Rosseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley and Sons, 1990.

P. F. Lazarsfeld and N.W. Henry. *Latent structure analysis*. Boston, MA: Houghton Mifflin, 1968.

S. Lê, J. Josse, and F. Husson. FactoMineR: an R package for multivariate analysis. *Journal of Statistical Software*, 25(1), 2008.

B. Le Roux and H. Rouanet. *Geometric data analysis, from correspondence analysis to structured data analysis*. Kluwer Academic Publishers, 2004.

D. A. Linzer and J. B. Lewis. poLCA: An R package for polytomous variable latent class analysis. *Journal of Statistical Software*, 42(10):1–29, 2011. URL `http://www.jstatsoft.org/v42/i10/`.

S. Liverani, D. I. Hastie, L. Azizi, M. Papathomas, and S. Richardson. PReMiuM: An R package for profile regression mixture models using dirichlet processes. *Journal of Statistical Software*, 64(7):1–30, 2015. URL `http://www.jstatsoft.org/v64/i07/`.

A. L. McCutcheon. *Latent class analysis*, volume 64 of *Latent Class Analysis*. Sage, 1987.

J. Molitor, M. Papathomas, M. Jerrett, and S. Richardson. Bayesian profile regression with an application to the National survey of children's health. *Biostatistics*, 11(3):484–498, 2010.

P. J. Rousseeuw. Sihouettes: a graphical aid to the interpretation and validation of cluster analysis. *Computational and Applied Mathematics*, 20:53–65, 1987.